

Bitcoin Forensics: A Tutorial

David Neilson, Sukhvinder Hara and Ian Mitchell

Middlesex University, London, UK

{d.neilson,s.hara,i.mitchell}@mdx.ac.uk

Abstract. Over the past eighteen months, the digital cryptocurrency Bitcoin has experienced significant growth in terms of usage and adoption. It has also been predicted that if this growth continues then it will become an increasingly useful tool for various illegal activities. Against this background, it seems safe to assume that students and professionals of digital forensics will require an understanding of the subject. New technologies are often a major challenge to the field of digital forensics due to the technical and legal challenges they introduce. This paper provides a set of tutorials for Bitcoin that allows for learners from both backgrounds to be taught how it operates, and how it may impact on their working practice. Earlier this year they were delivered to a cohort of third year undergraduates. To the author's knowledge, this represents the first integration of the topic into a digital forensics programme by a higher education provider.

Keywords: Bitcoin, Blockchain, curriculum design, digital forensics

1 Introduction

Bitcoin [1] is a decentralized cryptocurrency and payment network that allows for transactions to be conducted peer-to-peer amongst its users. Introduced in 2009, it has rapidly gained traction and currently has a market capitalization of almost ten billion dollars. A huge industry has grown to support these developments and the currency is now accepted as a method of payment by a number of large retailers. It has spawned a wide number of imitations and there are currently over two hundred cryptocurrencies which are referred to as alt-coins. There is now little doubt as to whether these technologies are here to stay and are likely to grow even further over the coming years. Bitcoin introduces a number of challenges to digital investigations mainly due to the elements of anonymity it provides a user, and the decentralized network which it operates within. These features have made it an attractive means of exchange for those engaged in criminal activities online. It has the potential to be used for money laundering and tax avoidance [2], and has been used extensively to purchase illicit goods and services through online marketplaces such as Silk Road [3]. It is also used as the method of payment for ransomware attacks which have witnessed huge growth during this time [4]. Currently there exists little in the literature to suggest that practitioners have the requisite knowledge of this technology to investigate crimes with Bitcoin. This situation is likely to change in the coming years and serves as the main motivation for the development of a training course in the form of seven tutorials.

Any learner, whether from Public/Private/Financial Sector, Security Agencies, Law Enforcement Agencies or Education, will rely on certain Learning Styles [5] with varying levels. In addition to learning styles, the framework also has to accommodate effective formative and summative feedback, to ascertain if learners have achieved learning objectives.

Often there is a reliance on experience when developing and designing something new, this applies to most areas, e.g. software engineering. Is it possible to prepare and design a curriculum without reliance and dependence on experience? LEAF [6] has been chosen partly answer this question and support curriculum development and learners' learning styles, and allow the integration of effective formative and summative feedback. The disadvantage of the framework is that it shifts the burden from the learner to the facilitator to provide good preparation that is shown from the study to be positively correlated with good results. Briefly, LEAF has two phases: i) design; and ii) implementation. In design phase, there are three stages: Lecture; Exercise; Apply and there are four stages in the implementation phase: Lecture; Exercise; Apply; and Feedback and hence the acronym. LEAF was chosen since previous results show that on average 60% achieve over 70% in assessment, this is important in the delivery of material to professionals, which often require a higher pass grade than Higher Education Providers, typically 40%, for accreditation or membership.

The approach suggested in this paper allows for the following tutorials to be applied by learners from both academic institutions in teaching of undergraduates, and also professionals already employed in the industry.

2 Background

Bitcoin is a distributed payment network that is built on the foundation of a number of key technologies; public key cryptography, p2p networks, cryptographic hash algorithms. These combine to produce a payment system that requires no central authority to validate transactions conducted amongst its peers. Fundamental to this system is a data structure called the blockchain, which is essentially a database containing every transaction that has ever taken place in this payment network. There are two main features of this dataset that makes it extremely secure as a store of data and therefore of use to digital investigators. Firstly, the data written to it is immutable, which is achieved through its use of a proof-of-work algorithm [7]. Once written to, it is practically impossible to modify the data contained within it. Secondly, a copy of it is maintained by every node within the Bitcoin network, meaning there is no single point of failure. These features make it an extremely useful source of evidence for the digital forensics practitioner.

However, the data contained within it is quite limited and the design of the system makes establishing the identity of its users extremely difficult providing a level of anonymity. There is no concept of accounts, and transactions are conducted between addresses that the user holds in what is called a wallet. Transactions are broadcast as messages to the network and propagated peer-to-peer, containing no IP address data

that would indicate where the transaction originated from. In counter to this, every transaction and its associated addresses are permanently recorded in the blockchain and viewable to anyone who has a copy. Each and every transaction is chained together, allowing for the path that funds take to be traced back to the first transaction. For these reasons Bitcoin is said to be pseudonymous rather than anonymous. Its current lack of acceptance as a major means of exchange means that at some point there will be conversion to Fiat currencies, and this provides investigators a means to reveal suspects identities, due to the regulatory requirements placed on money transmitters.

The lack of an established framework, and peer accepted methodologies to process bitcoin related evidence, may leave investigators without the required skills to conduct an investigation. For technically capable professionals, the learning of the subject should not pose too many problems. However an understanding of the general operating environment is also required, and in the event of an investigation this may be time that could be damaging to the pursuit of a suspect.

3 Course Design

3.1 Lecture, Exercise, Apply

The course is split into seven tutorials which are composed of both a lecture and an associated lab session. The lectures provide description and explanation of the key concepts and components, supporting theoretical and activist learning styles. This knowledge will then be applied through a number of exercises using Bitcoin, supporting pragmatic and reflectivist learning styles. The knowledge gained from each of the tutorials is directly relevant to parts of the assessment at the end and represents the apply stage of the LEAF model. The elements related to forensics were gradually introduced over the series of tutorials. This is so that learners can concentrate initially on understanding how the network and operating environment work, without any distraction from investigatory issues.

The lectures are designed to break down the main components of Bitcoin into easily digestible sections, which reflect the main topics of interest. The exercises were designed to follow on from the lecture and allows for learners to gradually familiarize themselves with the tools and software. The learners conduct a number of different transactions between their desktop wallets, and also a web based wallet they create in the lab. An important component of the exercises is that the learners are provided with a log sheet, which they use to record the pertinent details of every transaction they carry out. This forms an extremely important part of the curriculum design as they promote good practice of maintaining contemporaneous notes, which are then used by the learners to trace their own activity throughout the individual lessons. When the learners reach the later labs and start to analyze the blockchain, they will use these log sheets to verify the transactions they have carried out.

The apply stage comes at the end of the lecture and lab sessions, and is designed to test their understanding of the concepts and skills they have learnt. A crime scenario involving bitcoin transactions has been created that allow learners to explore a

number of pieces of evidence, which enables them to conduct an investigation. For the students we have taught at undergraduate level, they are asked to produce a report in the style of a digital investigation. This forces the learner to consider the investigation in the same manner as they would with a traditional one. Due to time considerations this method for assessment would probably not be appropriate for professionals, but could be replaced with a presentation of their findings and results.

3.2 Lab Setup

Good planning is essential to the successful delivery of the course and there are a number of things to organize and prepare. The first consideration should be given to which piece of wallet software should be used on the lab workstations, from where the learners will carry out most of their transactions. The ideal choice here would be to use a complete client such as Bitcoin Core. This is known as a full node and allows for all functions within the network to be carried out; transactions, mining, verification and propagation. To carry out these functions requires it to maintain a full copy of the blockchain which at the time of writing stands at 75 Gigabytes, and this creates maintenance issues. To keep the blockchain up to date the workstations must be left on throughout the teaching period, or they must be switched on periodically to let them catch up before the delivery of the next tutorial. The second option is to use a wallet that is called a Simplified Payment Verification (SPV) client, which does not keep a copy of the blockchain. While these types of client are slightly weaker from a security perspective [8], this will have no impact upon their use for the tutorials purposes.

A number of other pieces of software are also required to follow this tutorial plan. Numisight is a blockchain transaction explorer that generates a graph of transaction inputs and outputs, and is used in Tutorial 6. It is essential in helping learners to learn blockchain analysis techniques and also to identify the nature of a transaction structure. Due to the complexity of the data used, a visual representation of these connections can improve both the quality and speed at which this analysis can take place. Learners will also require access to an internet browser throughout all of the tutorials for two main purposes; to let learners create and use web based wallets to carry out transactions, and also to make use of online block explorers which learners will use to verify their transactions. The workstations should also have Python installed, so that they can utilize the script which simulates the process of mining used in Tutorial 4 The final requirement is that there needs to be access to the parts of the blockchain that contain the transactions they have carried out. If the Bitcoin Core client has been employed, then every workstation with the software will have a full copy of the blockchain. For those using SPV clients however it will be necessary to keep a record of which blocks are mined during the tutorials. This is used for the lab exercises where the learners manually parse transactions on the blockchain with use of a hex editor in Tutorial 5.

Learners must have enough bitcoins to be able to carry out transactions and pay the associated transaction fees, which represents the only direct financial cost for the delivery of the tutorials. Each cohort of twenty learners would require approximately

0.025 BTC which is worth just over fifteen US Dollars at the time of writing, and of which 0.005 would be reclaimed at the conclusion of the course. The final consideration is to have digital copies of the log sheet that is to accompany the lab exercises. While most practitioners still maintain handwritten contemporaneous notes, the extensive use of long cryptographic keys makes this approach impractical for this type of investigation.

3.3 Assessment

A crime scenario has been developed through the creation of a number of web based resources, and a series of Bitcoin transactions. This allows learners to explore all of the topics that have been looked at during the tutorials and apply them through a hypothetical investigation. It involved the creation of a number of accounts and leaving traces of activity so that the learners could build a case around the information they find. The assessment is designed to test the learners understanding of the Bitcoin environment and to display the techniques they have learnt over the period of the course. Learners are expected to analyze the transactions with use of the appropriate methods to trace and verify their findings, determine the leads generated by the evidence and to suggest what further steps could be taken to locate the suspect, and assess how much evidential weight they give to their discoveries and the limitations of them. The relevant information is not difficult to locate, as it was important to ensure that less capable learners are not disenfranchised from completing the assessment effectively.

The learners are informed that law enforcement arrived at the suspects home with a warrant to search her premises. Upon arrival they find that the suspect appears to have fled and there are no relevant possessions to be found in the home. They find one piece of evidence in her dustbin; a piece of paper that has three items handwritten on it; a Bitcoin address, the term 'Coingenie', and the address of a Tor hidden service indicated with the suffix '.onion'.

Bitcoin Address. The bitcoin address provided, is the address that is used to purchase the fake documents from the Tor website in the scenario. This places it at the center of a distinct chain and this allows the learners to try and establish the source and destinations of the bitcoins transferred. The transactions were designed so that it would be easy to pick up the trail, and make sense of the relevant chains to the case. Here the learners are able to detect some the transaction types and patterns that they have learnt. Critical to this, are the amounts that the transactions represent, as these can be linked to the values of some of the forged documents available for sale on the Tor website. Complicating this is the fact that when these funds are sent, transactions for almost exactly the same amount are moved from the receiving address to another, the monetary difference due to the transaction fee. This tests the student's ability to organize the correct transaction. The transactions are designed so that the issue of address reuse is apparent and the implications that can be taken from this activity. Depending on which blockchain explorer is employed to assist their investigation they

may manage to discover that the funds are transferred to a leading cryptocurrency exchange, and allows the learners to suggest that this may be in order to launder the bitcoins, a topic that is covered in Tutorial 7.

Coingenie. An account was setup on the popular bitcoin forum bitcointalk.org. Here the user had left a series of posts in two threads they had created, that were designed to reveal some information about Jean herself and some possible leads. The first thread was entitled ‘newbie needs a new wallet’. In this she states that she is looking for a wallet that is web based and in the next post indicates that she may use Circle, a well-known provider. This would allow learners to suggest that a request could be made to the service provider, to try and reveal the identity of the user. They learn in Tutorial 2 that Circle requires identity to sign up and buy and sell bitcoins, complying with anti-money laundering regulations. The other important thing that they could suggest is that the nature of her posts reveals that she has limited understanding of bitcoin and so is unlikely to adopt techniques such as changing addresses and not reusing. The second account created was at localbitcoins.com. This is a well-known method that allows users to meet and exchange bitcoins for cash, and is a means by which a user can mask their identity. An advert was setup to sell bitcoin from a location in England, the same country for where the transactions indicated the fake documents had been purchased for. An unexpected benefit of this term is that it produces a number of search engine results that were not created by the authors, allowing for the use of their logic to establish that they were not relevant.

Tor Hidden Service. This was used to setup a website and is designed to draw on what they had been taught through the case study of Silk Road in Tutorial 7. When accessed using Tor, learners are taken to a webpage that advertises the selling of fake documents. There are a few important pieces of information that the student may glean from here, the most critical being the values of the documents that are on sale. Passports, ID cards and Drivers licenses are available for a number of different European countries and importantly, different values are given for each of them. This enables the learners to find a link between these and some of the transactions that linked to the bitcoin address they are provided with. The suggestion made by the transaction data suggests that two UK documents have been purchased. This also allows a link to be made with the localbitcoins.com account which was also based in the UK. The website’s FAQ also contains a subtle clue, which helps to suggest without any doubt which transactions were responsible for purchasing the documents. The transaction chain has an almost identical transaction value right after the documents have been purchased, the difference being the transaction fee. The FAQ states that customers should “not ask for discounts” and this suggests that learners should look very closely to determine the correct transactions.

4 Tutorials

4.1 Tutorial 1

Lecture. The first lecture provides an overview of the entire system so that learners can appreciate what it is, and how it works at a simple level. This provides the foundation for the later lectures, where the significant components are dealt with in more detail. The lecture begins by starting with a short history about what money is. This is important as it allows the student to appreciate that there have been a variety of methods used as a means of exchange, allowing them to appreciate the similarities and difference with cryptocurrencies. Next the main technologies that Bitcoin utilizes are described; public key cryptography, peer-to-peer networks, and cryptographic hash functions. It is essential that a student has at least a basic understanding of these before trying to look at the topics in greater detail, so they can appreciate the design aspects that make Bitcoin a secure system. Each of the main components are explained, and the overriding objective is to get learners to understand that transactions are conducted peer-to-peer without the need for a central authority to verify them. The session concludes by showing the example of a transaction and all of the main steps that are involved in processing one.

Exercise. The first lab is mostly concerned with setting up all of the relevant accounts and wallets that the learners will use to carry out transactions in the subsequent labs. They are also given a copy of the log sheet and record details of their accounts into these. The learners all send a copy of a receive address to the teacher so that they can distribute the bitcoins that they will be using for all of the tutorials exercise. The main objectives for the exercise session are for the student to become familiar with the wallet software that will be used, and to conduct their first transactions.

Apply. Due to the simple level with which the material is dealt with at this point, there is not too much here that helps to inform the student greater about their coursework. One part that is perhaps most important here is that the units and notation of bitcoin is introduced. This can have implications for their work in the labs, and also on the correct presentation of their work for their assessment.

4.2 Tutorial 2

Lecture. The two main topics for the second tutorial are the storage and acquisition of bitcoins. The most important concept for the learners to understand here is that wallets do not store bitcoins, they store addresses. The first part of the lecture describes the process for creating addresses and this necessitates a need to explain how public key cryptography works. This is done at a relatively high level, as it is not necessary to have a deep understanding of the ECRSA algorithm used by Bitcoin. They are then shown how the bitcoin address is converted from its public key form through hashing and finally encoded as a Base58 string. The next part of the lecture concentrates on how a user is able to purchase or acquire bitcoins. The main methods are discussed

and allows for a comparison between their relative merits and disadvantages. Of most importance here are the levels of anonymity and the possibility for a criminal to prevent tracing. The potential to use the system for money laundering is introduced, and the point is made that the weakest point in the system is the point of exchange with the current financial system. Finally, the lecture looks at wallets and the various different methods of storage users have at their disposal.

Exercise. During Tutorial 1's lab session, the transactions the learners make use the default addresses provided by the wallets they create. In the first exercise for Tutorial 2, the learners conduct transactions by creating and making use of new addresses in both their desktop client and also their web based wallets. They also specify new change addresses for these transactions to reinforce the concept of avoiding reuse of them. The second exercise requires the learners to work in small groups and asks them to sign up for some other online wallet providers, allowing them to compare the different security requirements for signup and also the security features they provide to the user. The final exercise involves the learners creating backups of their desktop wallets and also the creation of paper wallets to securely store a very small amount of bitcoins from their allocation.

Apply. One transaction within the crime scenario shows quite explicitly that an address has been reused, and allows the student to conclude that the purchaser of the first fake document was the same person who purchased the second document. The learners are also given information in a bitcoin forum that the suspect is considering using an online service provider. This allows the student to put in their report what actions would be helpful given this information. They should be able to say that they would contact the service provider who may provide either an IP address, or the name and address details linked to a payment card.

4.3 Tutorial 3

Lecture. Tutorial 3 looks at transactions in detail and the learners are first given a quick overview of addresses as they are the main components of the transaction. They are shown how each transaction is broadcast as a message to the network, which is then propagated to each and every node for validation. The main part of the lecture focuses on the construction of these messages, and how transactions consist of a number of input and output addresses [9]. These all form a chain of transactions that is important in terms of tracing the source and destination of funds. The most common transaction forms and types are described showing how it is possible to have multiple inputs and outputs for any given transaction, and how in most cases the need for a change address will mean that every transaction will usually have at least two outputs. A byte descriptor table is introduced (Table 1 below) showing the specific details of the data structure, and help explain the interaction between private and public keys, through the use of locking and unlocking scripts.

Table 1. Byte offset descriptions of transaction structure

Field	Size	Description
4 bytes	Version	Specifies which rules this transaction follows
1–9 bytes (VarInt)	Input Counter	How many inputs are included
Variable	Inputs	One or more transaction inputs
1–9 bytes (VarInt)	Output Counter	How many outputs are included
Variable	Outputs	One or more transaction outputs
4 bytes	Locktime	A Unix timestamp or block number

Exercise. The exercises for Tutorial 3 consist of carrying out a number of transactions between their desktop and web based wallets. They also explore the transaction functionality of the wallets such as the ability to define the size of transaction fee and the use of labelling. These are done with specific amounts so that different types of transaction are simulated, which will then be analyzed when the learners revisit them in Tutorial 7’s lab session. A transaction is carried out with the same change address as the input address so that when the student analyses them later this action of reuse will be very apparent to them. For the second part of this lab the learners are placed into groups to research and make a comparison of different web based block explorers. Using some of their previous transactions they are asked to determine which one provides the most useful features and then discuss with the rest of the cohort.

Apply. The assessment requires that the student understand a number of concepts and resources that they will use in their investigation. The chaining of transactions and how they can be traced with block explorers, provide the first introduction to blockchain analysis. Understanding how transaction fees are calculated as implied as part of the transaction and not included as an output, can be critical to the learners correctly identifying which transactions were used to purchase the fake documents in the assessment scenario. The issue of address reuse is also covered again here.

4.4 Tutorial 4

Lecture. Tutorial 4 looks at the peer-to-peer network, and the components that allow it to achieve consensus in the absence of a central authority. The Bitcoin protocol and the main types of nodes are explained, and how all these parts are required to achieve

a shared view of the networks state. During the course of the lecture the learners are introduced to what Antonopoulos refers to as the ‘four pillars of consensus’ [10]; verification of every transaction by each node in the network during propagation, collection of transactions into blocks and processed through use of a proof-of-work algorithm, verification of newly created blocks by each network node and added to the blockchain, each nodes selection of the longest chain of blocks. The process of mining is explained and how through their use of a proof-of-work algorithm, they secure the network against subversion and attack. The process of new coin generation is described, and how this provides incentive for the miners to commit their hashing power. The overall aim for the lecture is to allow learners to appreciate the system as a whole.

Exercise. Now that the learners have seen all of the components work they should be able to understand how they all work together. However due to the systems complexity this may not be the case, and it is vital that all the learners understand this before progressing to forensic analysis. This led to the design of an exercise whereby the learners become components of the network, to carry out a physical simulation of a transaction being processed from broadcast to its commitment to the blockchain. Essential to this is that two learners act as miners competing to find a hash below a certain value so that they can claim the coinbase reward for doing so. The following Python Script has been developed for this purpose;

```
import hashlib
text = ("Miner 1 this is another example of the hash of a
block header")
for nonce in range(10000):
    input = text+str(nonce)
    hash = hashlib.sha256(input.encode()).hexdigest()
    print (input, hash)
    if hash.startswith("00"):
        print ("Found Hash")
        break
```

The value found within the *if* statement of the code (representing the nonce) can be modified to search for more zeroes which will let the program run for longer before finding a match. This enables the concepts of hashing power and hash difficulty to be demonstrated very clearly. For the second miner involved in the demonstration, modifying the variable text to “Miner 2” will result in different hashes being generated for each miner

Apply. While there are no distinct parts to this tutorial that are directly related to their assessment it does represent one of the most fundamental parts to the course. An understanding of this tutorial is essential for the learners to make sense of the investigation as a whole.

4.5 Tutorial 5.

Lecture. Up to this point the blockchain has not been looked at in any detail and has simply been treated as a database that keeps a record of every transaction, and therefore can validate whether a transaction can be processed. In this tutorial the data structure is looked at in detail and is broken down with use of the byte tables describing the various fields. The learners are shown how the hash of the previous block header field, is the essential piece of data in maintaining the chain of blocks, and prevents modification of the data within the blockchain. The two methods for identifying a block are described; the block hash (a hash of the block header) and the block height, with the latter of these not being a unique identifier due to the possibility of forks being generated in the blockchain. The potential for two copies of the blockchain to temporarily coexist is explained, as is the way in which these forks are resolved. The lecture concludes by describing the various types of chain, and their relationship with the ‘longest chain’, which all nodes accept as being the correct version of the blockchain.

Exercise. The learners are provided with a block.dat blockchain file and are taken through an example by the teacher with each step clearly explained. Each field is located with a hex editor and a description of the various fields shown in Tables 2 & 3 below. These are then converted to determine the values and compare them with the results of an online block explorer. The learners are then asked to do the same for the first transaction they have carried out in Tutorial 1’s lab session. The relevant block will need to be provided, and due to the fact that they were all created at the same time, all of the cohort’s first transactions should be within the same blockchain file. Due to the presence of variable length fields, this task can be fairly demanding on the learner and the teacher may need to repeat their example again.

Table 2. Byte offset descriptions of the block

Field	Size	Description
Block Size	4 bytes	Size of block in bytes
Block Header	80 bytes	For details view Block Header (Table 2) below
Transaction Counter	1-9 bytes (VarInt)	How many transactions in block
Transactions	Variable	The transactions (see transaction byte table)

Table 3. Byte offset descriptions of the block header

Field	Size	Description
Version	4 bytes	Version No. to track protocol upgrades
Previous Block Hash	32 bytes	Reference to hash of parent block

Merkle Root	32 bytes	Hash of the root of the Merkle tree for this block
Timestamp	4 bytes	Approximate block creation time (Unix Time)
Difficulty Target	4 bytes	Difficulty of POW algorithm for this block
Nonce	4 bytes	Counter used to generate POW algorithm

Apply. As part of the assessment learners are expected to verify the transactions. This can be carried out with use of online block explorers, but this data should also be verified manually. The learners are provided with the block.dat files from the blockchain that allow for manual parsing. This is quite a difficult task as the only way in which they are able to locate the relevant transaction is by making use of the Bitcoin API function *addresstohash*. This allows for the address used to be located and represents the first step in finding the data in the blockchain. The other parts of this transaction can then be verified through use of the field descriptions shown in Table 3.

4.6 Tutorial 6.

Lecture. The main topic for Tutorial 6 is crime that involves the use of bitcoins. Before this however the potential for the network to be compromised is considered, beginning with a detailed look at how forks within the network develop and how they are resolved. The learners are reintroduced to the concept of a 51% attack, and presented with a hypothetical case showing how it could be carried out. Other attacks such as denial of service are also considered as is the ability to detect the IP address of the user who makes a transaction [11], although it is made clear the limitations of these methods. The second half of the lecture looks at the types of crime that have been carried out using bitcoin and the methods they employ to remain undetected. This begins with a reminder of what makes it attractive to criminals, looking in particular at the huge rise in ransomware attacks. The collapse of former bitcoin exchange Mt Gox is also examined, and for both of these cases, the ability to trace the proceeds of crime through the blockchain. The infamous case of the Silk Road marketplace accessed through Tor is described, and how the funds could be linked to the suspect due to his laptop containing a wallet associated with the relevant addresses. The final part of the lecture looks at the methods that can be employed to prevent tracing such as coinjoin transactions and bitcoin mixing services.

Exercise. In this lab session the learners conduct their final transactions for the course, all of which are analyzed in the final lab session. They are asked to send all of their remaining bitcoins to an address setup by the teacher on a different wallet to the one they use to distribute the learner's bitcoins in Tutorial 1. The teacher then sends to

funds to bitmixer.io, a bitcoin mixing service that takes incoming bitcoins and breaks the link with. This allows for the transaction to be analyzed in the final tutorials exercises. While it would be advantageous for learners to do this by themselves the min transfer amount of 0.01 BTC places pressure on the cost of the courses delivery. Due to their own chains being linked to this the learners can still look at this data when they start to analyze their own history in Tutorial 7. The class is also introduced to Tor, and they are asked to explore the environment to understand the scale and variety of criminal behaviour that is occurring. The learners then share their experience and findings with the rest of the cohort provoking discussion around the implications.

Apply. The tutorial has provided details of the types of crime involved and this provides some hints as to what may have happened in the crime scenario they have been studying provides them with the knowledge...The coin mixing exercise has particular relevance to the assessment. The crime scenario's chain of transactions eventually leads to an address that suggests the intention is to launder the proceeds of the crime. The use of Tor will also form part of their assessment, and this part of the tutorial provides them with knowledge of how it works, and the limitations this also places on an investigation.

4.7 Tutorial 7.

Lecture. The final lecture deals with blockchain analysis and how to make sense of transaction history and chaining. The learners are reminded of the benefits that the blockchain provides investigators, in particular the transparency and auditability it provides. The design of the system means that transactions effectively form a directed acyclic graph, and learners are exposed to some of the academic work that has been carried out in this area [12] [13]. The need for visualization is stressed and also the benefits this can bring in terms of making the evidence easier for a jury to understand. The lecture integrates a discussion with the learners as to what data should be displayed, and also the most effective way in which to do so. The learners are then introduced to Numisight, a blockchain explorer which as shown in Figure 1 displays the transactions as a graph, and enables manipulation of different display variables. This style is then used to describe the most common patterns of transaction that are found, and the types of behavior these may represent.

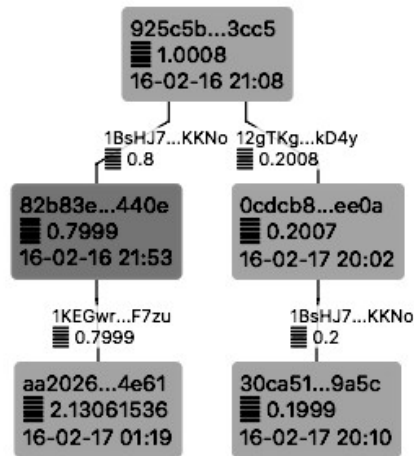


Fig. 1. Numisight blockchain explorer used to produce graphs of transactions.

Exercise. Having conducted a number of transactions of different types throughout all of the lab sessions the learners have created their own chain of address and transaction use. In this final lab session, the learners make use of this to perform the types of analysis they learnt about in the proceeding lecture, using their own transaction history. This is an important element of the courses design, and has the benefit of making it easier for the learners to identify with and understand the transactions, which is helpful considering the relative difficulty of recognizing long strings of random characters used in a lot of Bitcoins data structures.

Apply. This final tutorial provides the awareness and skills needed to effectively describe the transactions that take place within the assessment. The transaction types they have learnt about are much easier to detect visually and some of the ones described in this tutorial are very apparent in the crime scenario, such as relay and self-spending. It also makes the learners aware of the different ways in which it can be displayed allowing for the learner to develop their own methods for presenting the evidence.

5 Delivery

The main aim and motivation of this research is to develop a clear and concise curriculum on Bitcoin that can be delivered to students and professionals alike. Typical delivery in a Higher Education Institute/Provider to a student base typically is over 6 weeks, composing of 1 hour lecture, 2 hour lab/tutorial and a further 3-4 hours

independent study with suggested reading. An optional further two weeks can be added for summative and formative feedback. Such a delivery would require that prerequisite learning outcomes have been achieved elsewhere on the course, e.g. Encryption, and if not would extend the delivery by a further week. With this in mind it is possible to see that by the time you allow for effective formative feedback (assuming an in-course assessment is set) then this could easily extend to 12 weeks, or a term. Typical delivery to professionals is over 3 days; with a lecture and tutorial in morning (9:30-13:00) and afternoon sessions (14.00-16:30).

Reflective and Activist learning styles are important to the understanding of a subject, but often require total understanding of the fundamentals in technology-based fields, e.g. a basic understanding of Bitcoin security is required before it can be challenged. Generally, there may be exceptions, LEAF covers learning styles as follows:

- Lecture – Theoretic and Activist
- Exercise – Pragmatic and Theoretic
- Apply – Pragmatic and Theoretic
- Feedback – Activist and Reflective

In the design above some prudent managers may expect the tutorial to be squeezed into 2 days, however, this would not allow questions and answers sessions that have been designed to extend and test learners understanding at the end of each tutorial. This promotes feedback to the learner and accommodates reflective and activist learning styles.

The assessment scenario that has been developed allows for learners to be tested in a number of formats. For undergraduates this ideally should be the production of a forensic report, which they should have had some experience of during their degree studies. The amount of material that can be developed lends itself very well to students working in teams of three or four students. However for professionals who may well have been covering the material within an intensive short period this is not appropriate, and can easily be adapted to working in teams and then providing a presentation of their findings. Their existing skillset means this is likely to be achievable.

6 Conclusion

The main contribution of this paper is to satisfy the aim of developing a clear and concise curriculum, using LEAF, for a Bitcoin tutorial. Many tutorials are difficult to get right the first time of delivery and experience is relied on during the design phase or a lack of preparation shifts the burden onto the learner and yields the statement, “the easiest way to teach leads to the hardest way to learn”. Here the LEAF framework has helped get this balance correct. However, LEAF goes on to show that good preparation are related to good results, and therefore the shift of burden yields the converse of the previous statement, which is, “Easiest way to learn is the hardest way to teach”.

The other major contribution this paper makes is in proposing a unique method of assessment that is strongly linked to professional practice. While hypothetical in nature it provides an environment in which learners can make use and develop their skills for a new area that is still to be defined. While the paper is not trying to develop a framework for Bitcoin investigations, it does suggest a number of approaches that may be applicable to future investigations. It also helps learners to think beyond the restrictions of what they have been taught, and this type of adaptability maybe useful if other new technologies arrive which may also bring challenges to the field of digital forensics.

At the time of writing the results for the first students to have participated in the course have just been received. Of the thirty-six students who took part, only 3 of them failed to achieve a pass grade of 40%. This represents a success rate exceeding 90% of the students and endorses the use of LEAF to help in the first delivery of a new topic for digital forensics.

References

1. Nakamoto, Satoshi. "Bitcoin: A peer-to-peer electronic cash system." (2008).
2. Möser, Malte, Rainer Böhme, and Dominic Breuker. "An inquiry into money laundering tools in the Bitcoin ecosystem." In eCrime Researchers Summit (eCRS), 2013, pp. 1-14. IEEE, (2013)
3. Ron, Dorit, and Adi Shamir. "Quantitative analysis of the full bitcoin transaction graph." In International Conference on Financial Cryptography and Data Security, pp. 6-24. Springer Berlin Heidelberg, (2013).
4. Kharraz, Amin, William Robertson, Davide Balzarotti, Leyla Bilge, and Engin Kirda. "Cutting the gordian knot: a look under the hood of ransomware attacks." In International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment, pp. 3-24. Springer International Publishing, (2015).
5. Honey, Peter, and Alan Mumford. "The manual of learning styles." (1992).
6. I. Mitchell and M. Sheriff. "Lecture, exercise, apply and feedback: Turning a new leaf in curriculum design". In 18th Int. Conf. on Software Process Improvement, Research and Education, INSPIRE 2013. The British Computer Society (2013)
7. Back, Adam. "Hashcash-a denial of service counter-measure." (2002).
8. Biryukov, Alex, and Ivan Pustogarov. "Bitcoin over Tor isn't a good idea." In 2015 IEEE Symposium on Security and Privacy, pp. 122-134. IEEE, (2015).
9. Androulaki, Elli, Ghassan O. Karame, Marc Roeschlin, Tobias Scherer, and Srdjan Capkun. "Evaluating user privacy in bitcoin." In International Conference on Financial Cryptography and Data Security, pp. 34-51. Springer Berlin Heidelberg, 2013.
10. Antonopoulos, Andreas M. "Mastering Bitcoin: unlocking digital cryptocurrencies." O'Reilly Media Inc, (2014).
11. Koshy, Philip, Diana Koshy, and Patrick McDaniel. "An analysis of anonymity in bitcoin using p2p network traffic." In International Conference on Financial Cryptography and Data Security, pp. 469-485. Springer Berlin Heidelberg, (2014).

12. Meiklejohn, Sarah, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage. "A fistful of bitcoins: characterizing payments among men with no names." In Proceedings of the 2013 conference on Internet measurement conference, pp. 127-140. ACM, (2013).
13. Reid, Fergal, and Martin Harrigan. "An analysis of anonymity in the bitcoin system." In Security and privacy in social networks, pp. 197-223. Springer New York, (2013).